AD-A040 300 DECISIONS AND DESIGNS INC MCLEAN VA F/G 12/2

THREE PAPERS ON THE VALUATION OF DECISION ANALYSIS.(U)

MAY 77 S R WATSON, R V BROWN, D V LINDLEY N00014-75-C-0426

UNCLASSIFIED TR-77-2 NL

| OF |

AD
A040 300

END

DATE
FILMED
6-77

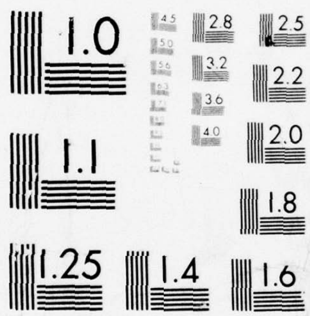

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

(12)

Three Papers on the Valuation of Decision Analysis

S. R. Watson
R. V. Brown
D. V. Lindley

DECISIONS and DESIGNS, INC.

Office of Naval Research
Contract Number N00014-75-C-0426
Contract Authority Identification
NR 197-029 (455)

TECHNICAL REPORT 77-2

THREE PAPERS ON THE VALUATION OF DECISION ANALYSIS

by

S.R. Watson
Engineering Department
University of Cambridge

R.V. Brown
Decisions and Designs, Incorporated

D.V. Lindley
University College, London

Sponsored by

Office of Naval Research
Engineering Psychology Programs
Contract Number N00014-75-C-0426

May 1977

DECISIONS AND DESIGNS, INC.
Suite 600, 8400 Westpark Drive
McLean, Virginia 22101
(703) 821-2828

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER (14) TR-77-2	2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) (6) THREE PAPERS ON THE VALUATION OF DECISION ANALYSIS.	5. TYPE OF REPORT & PERIOD COVERED (9) Technical Report.
	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) (10) S. R./Watson (University of Cambridge) R. V./Brown D. V./Lindley (University College London)	8. CONTRACT OR GRANT NUMBER(s) (15) N00014-75-C-0426
9. PERFORMING ORGANIZATION NAME AND ADDRESS Decisions and Designs, Inc. 8400 Westpark Drive McLean, Virginia 22101	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research 800 North Quincy Street Arlington, Virginia 22217	12. REPORT DATE (11) May, 1977
	13. NUMBER OF PAGES 48
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office) (12) 46p.	15. SECURITY CLASS. (of this report) UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

decision analysis
utility theory
resource allocation
rationality

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report pursues an ongoing investigation into techniques for assigning numerical value to proposed applications of decision analysis, in collaboration with faculty from Cambridge University and University College London. Three specific papers are presented related to:

The conceptual mathematical and graphical explication of some basic procedures; → next page

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

390664

LB

cont

An adaptation of the procedures to a real decision analysis, and

The notion of perfect rationality as a reference point against which proposed analyses can be evaluated.

SUMMARY

This report pursues an ongoing investigation into techniques for assigning numerical value to proposed applications of decision analysis, in collaboration with faculty from Cambridge University and University College London. Three specific papers about the conceptual mathematical and graphical explication of some basic procedures, an adaptation of the procedures to a real decision analysis, and the notion of perfect rationality as a reference point against which proposed analyses can be evaluated are presented.

1. <u>The Valuation of Decision Analysis</u> by S.R. Watson and R.V. Brown

The most obvious motivation for using decision analysis (applied decision theory), to offset its cost, is to improve the quality of the subject's decision. This paper explores approaches for quantifying this prospective value for a proposed analytic effort, on the assumption that, like unaided decision making, it will fall short of perfect rationality. One valuation approach compares the subject's expected utility with and without the proposed analysis. Its formal properties are examined, and a graphical implementation procedure is suggested. A second approach assesses the reduction in the expected cost of irrationality, a concept analogous to opportunity loss in the valuation of information.

2. <u>A Case Study in the Valuation of Decision Analysis</u> by S.R. Watson

This paper approaches the valuation of analysis derived in (3) to a case study with which the author is involved.

We deduce several numerical measures for the value of a specific analytical exercise which differ considerably but are of the same order of magnitude. We conclude that the framework of (3) is a useful means to discuss the amount of effort that it is rational to spend on analysis in order to test its operational feasibility, and we suggest promising directions for further development. This initial case study was taken from the commercial sector in order to have as simple as possible a measure of the value financial profit. This choice was to pave the way for the more complex multi-attributed situations commonly encountered in defense and government.

3. <u>Notes on Rationality and the Reconciliation of Incoherence</u> by R.V. Brown and D.V. Lindley

The principle that coherence between a subject's assessments is to be sought in decision analysis is well established.

This paper explores additional principles for reconciling a subject's initial incoherence with a view to defining a "most rational" analysis against whose value any proposed decision analysis can be compared.

CONTENTS

FIGURES

PREFACE

The task of developing a methodology to assign numerical values to the benefits accruing from proposed applications of decision-analytic techniques was initiated in 1975 as part of the Office of Naval Research's ongoing program in decision analysis research, monitored by Dr. Martin A. Tolcott. The principal investigator for this task has been Dr. S. R. Watson of Cambridge University under the technical direction of Dr. R. V. Brown of Decisions and Designs, Incorporated and University College, London. Dr. D. V. Lindley of University College, London contributed time to the elucidation of key ideas related to the definition of rationality and the resolution of incoherence.

This report consists of three separate reports bearing on different aspects of the total inquiry.

The first report, "The Valuation of Decision Analysis," establishes a conceptual and computational framework for the project as a whole. It is an extension and refinement of material previously presented in a tentative form in an earlier DDI report, <u>Issues in the Value of Decision Analysis</u>, Technical Report 75-9, by S. R. Watson and R. V. Brown. The present report has been accepted by <u>The Journal of the Royal Statistical Society</u> for publication in Series A.

The second report illustrating one specific approach in use is based on a case study developed for and in conjunction with a major British company, and the material is cited with their permission.

The third report is a summary of a working paper which emerged from theoretical discussions between Dr. R.V. Brown and Professor D.V. Lindley at University College, London in the fall of 1976. This material is being developed with the collaboration of Dr. Amos Tversky under a current ONR contract.

ACKNOWLEDGMENTS

We would like to express our thanks to Professor D. V. Lindley, Dr. M. A. Tolcott, Dr. T. W. Keelin, III, and Dr. D. E. Bell, all of whom made helpful comments on earlier drafts of this paper.

1.0 THE VALUATION OF DECISION ANALYSIS

by S.R. Watson and R.V. Brown

1.1 Introduction

Over the last thirty years there has been a great deal of effort devoted to developing quantitative techniques to help people make decisions. Proponents of these techniques usually assert that their use improves the quality of decision making; by this they mean that the goals of the decision maker are more nearly fulfilled by their use than by other approaches to making decisions. In some circumstances it is clear, once the problem has been analyzed, that this is indeed the case. The use of linear programming in a trim-loss problem, for example, produces a solution that cannot be bettered. In the presence of uncertainty, on the other hand, there is greater room for controversy. Suppose a decision maker decides to use the techniques of decision analysis (see, for example, Raiffa (1968), Lindley (1971), Brown et al. (1974)). If, in retrospect, it can be seen that a different decision would have resulted in a better outcome, the cynic can always claim that this particular technique was valueless or even dangerous; on the other hand, the believer can always argue that luck was against him. It is very difficult to argue for an objective test of the validity of a particular technique for decision making in the face of uncertainty. The most obvious candidate for such a test would be that, in the long run, better results were obtained by using the technique than by not so doing; but we doubt that a convincing controlled experiment to test this hypothesis could be constructed.

So far in this argument, we have made no mention of the cost of analysis. Nowadays it is quite common to spend large amounts of money on analysts and computer-time to apply some quantitative technique to a particular decision. Before doing an analysis, the responsible decision maker will want to be sure that he saves more by using the solutions provided by the new technique than he spends in using the technique itself. Even in something as clear-cut as using linear programming to schedule refinery operations, it is not entirely obvious that the computer and its programmers will pay for themselves prior to the establishment of a computer-aided scheduling system. This problem is even more acute in the use of decision analysis. If there is no objective measure of the value of decision analysis even after the analysis is completed, there is, a fortiori, no objective measure of its value prior to the expenditure of resources on it. This fact does not seem to prevent many decision makers from spending considerable sums of money on using this technique; in America, for example, hundreds of

thousands of dollars have been spent on the analysis of single decisions. One must presume that decision makers see some value in analysis. In this paper we argue that the framework of decision analysis itself provides a rationale for spending money on decision analysis, and we use that framework to establish a rational prescription to guide the decision maker who wants to know just how much money to spend on a projected analysis. In the next section we discuss the problem of value, and in section 1.3 we describe distinct approaches to its assessment. Section 1.4 contains details of a possible implementation of one of these approaches, and in the last section we discuss the difficulties that lie in the way of producing a practical tool to judge the value of decision analysis prior to its use.

1.2 A Rational Structure for the Valuation of Decision Analysis

It is reasonably uncontroversial to assert that decision making under uncertainty can be construed as a choice between gambles. Whichever alternative action is in fact taken, there is a measure of uncertainty about the final outcome. Of course, the outcomes from some alternatives may be certain, in which case we are stretching the use of language to refer to selecting one of those alternatives as "taking a gamble." However, since the problem is not that of decision in the face of uncertainty unless at least one of the alternative courses of action gives rise to an uncertain outcome, we lose no generality by representing the problem as a choice among gambles. The argument which supports decision analysis asserts that for each gamble the decision maker has a "certainty equivalent;" he is indifferent between taking the certainty equivalent as a sure outcome, and taking the gamble. The rational decision maker selects that gamble for which this certainty equivalent is most preferred. Thus, decision making in the face of uncertainty reduces to a choice among sure outcomes. It may be that good intuitive decision makers can always spot the gamble they prefer, but many people faced with a complex decision involving great uncertainties and important issues feel an alarming inability to tie all the individual judgments together. The technique of decision analysis decomposes the problem of comparison into many judgments that are much easier to make, and it ensures that an analysis of a problem makes the judgments consistent with one another. Independent assessments of the utility of the outcomes and of the probabilities of the uncertain events are made and combined to form the expected utility of each gamble. Then the previous judgments about utility can be used to deduce the certainty equivalents of each gamble. This step is not generally necessary, however, since utility is so defined that a sequence of increasingly desirable outcomes has increasing utilities. The rational decision maker merely chooses the gamble with the largest expected utility. If the deduced certainty equivalent of a

gamble does not correspond with an intuitive judgment of its value, then further analysis and adjustment is carried out until all the judgments cohere.

Many objections can be made to this framework as a practical prescription for decision making, not the least of which are doubts that adequate measures exist for determining degrees of belief in uncertain propositions or the value of outcomes. These issues have been recently discussed in this journal at some length (Hampton et al. (1973), Hull et al. (1973)). However, if it is desired to find some technique for decomposing a complex decision problem into simpler judgments, it is difficult to conceive of any procedure more compelling than that outlined above.

We shall go no further here in examining the philosophical basis for decision analysis although we recognize that it is far from satisfactory (see Watson & Brown (1975a Para. 3.0) for further discussion of this point). Rather, we shall use its normative framework, outlined above, to discuss the value of a proposed analysis. The presumption is that the decision maker has a set of subjective probabilities about the uncertain events facing him and a set of relative values about the outcomes which together logically determine which gamble he should take. The good intuitive decision maker is able to make this determination by a little introspection; the mechanisms of applied decision analysis are there to help the man whose intuition is not so sound. But whatever his intuitive abilities, there exists, for each decision maker and each problem, a von Neumann-Morgenstern utility scale describing his valuation of the outcomes and his attitude to risk, unique up to a positive linear transformation. We represent the expected utility of the most preferred option on this scale by U_P. We use the suffix P to indicate that this is the expected utility that the decision maker will attain if he analyzes his beliefs and values perfectly, and then takes the action with the largest expected utility. (For a discussion of whether it is meaningful to talk of perfect analysis, see Brown and Lindley (1976)). The value of analysis is related both to U_P and to the utility obtained if, where no analysis is done, the decision is made by some other means. Since this unaided decision procedure is itself a gamble, it must have an expected utility; let this be U_O, measured on the same scale as U_P. If now we conceive of paying a fee for the analysis, U_P will be reduced, since all the outcomes of the decision problem are less desirable (assuming only that the decision maker is not so perverse as to have a utility for money that strictly decreases at any point). The fee which reduces U_P to U_O is defined to be the value of analysis. It will clearly not be sensible to spend more than this on carrying out the analysis.

What we have defined so far is the value of perfect analysis; that is, the advantage the decision maker gains from uncovering the implications of all his beliefs and judgments. Every practicing decision analyst knows, on the contrary, that no analysis he carries out can possibly be considered perfect. Probability judgments, for example, are never likely to be so accurate that more time spent on elicitation could not improve them. Thus, when the decision maker is contemplating spending money on analysis, he has to recognize the extent of its imperfection. Since he is not likely to be sure in just what way the analysis will fail to be perfect, further uncertainty is introduced. But he is still facing an uncertain prospect; being a rational man, he has a utility for U_I for the certainty equivalent of this prospect. Thus, the value of the projected imperfect analysis is that fee which when deducted from all projected outcomes of the decision reduces U_I to U_O. For ease of explanation in what follows, we shall assume that the utility for money is linear, so that the value of perfect analysis is U_P-U_O and that of imperfect analysis is U_I-U_O; the reader will be aware that our arguments in no way depend on this assumption and are easily extended to the more general case.

There is obvious parallel here with the role of information in decision analysis. The value of perfect information is the difference between the expected utilities of two uncertain prospects, making the decision in the light of current uncertainties and resolving the uncertainties before the decision is made but after the fee is paid for the extra information. A similar comparison holds between imperfect analysis and imperfect information. Indeed, some previous discussions of this topic (Matheson (1968), Rice (1974)) have rested on this analogy. While the two concepts have appealing similarities, we ought to point out important differences. Analysis is concerned with manipulating the data known, by the decision maker, while information introduces new data into the problem. We can say that information reduces uncertainty, but analysis reduces confusion. Second, the expected value of information whether perfect or imperfect, can never be negative, so long as it is obtained at no cost. (See, for example, the arguments at the end of Chapter 7 of Lindley (1971).) On the other hand, it is possible for analysis to have a negative expected value. This apparently paradoxical assertion is perhaps best understood by reference to the framework for evaluating analysis which we develop in the next section; indeed, we give an example of this property towards the end of section 1.3.1. It arises from the fact that analysis can change the decision maker's preference ordering over alternative actions, not necessarily for the better. We generally suppose that the more analysis we carry out, the more nearly do we uncover the correct ordering, namely, that to be obtained from a perfect analysis. There is unfortunately, no guarantee that this is the case, since

it is conceivable that parts of the analysis could be so badly performed that the results of the analysis would be misleading. For example, suppose a decision maker has an initial belief that Action A is better than action B, but, on submitting to an imperfect analysis it turns out that B is better than A. If after impeccable further analysis this conclusion is reversed, so that perfect analysis reveals that A is indeed better than B, then the decision maker has lost utility if he acts on the basis of the imperfect analysis. That such reversals are possible could lead a decision maker to expect that, in some circumstances, a particular projected analysis might have negative value. Of course, if it did have negative value in prospect, he would disregard its findings. But the possibility of a negative expected value for a projected analysis exists, in contrast to the expected value of information, which, no matter how misleading the information, cannot be negative.

Our task now is to measure, for a particular decision and a particular decision maker, the value of the analysis, necessarily imperfect, that he contemplates, namely $U_I - U_O$. But before suggesting possible devices for doing this, we should point out that decision analysis may be used in practice for reasons other than to discover the most preferred gamble, as outlined above. We could say that this latter activity leads to the direct value of analysis, but it also has an indirect value. If a decision maker finds that the analytic approach reduces confusion in a particular case, he may well learn from this experience and improve his decision making thereafter. Moreover, when decision analysis is used within an organization, the members of that organization may find that communication of ideas and opinion among themselves is considerably enhanced by having the decision described by a model; this discovery could conceivably lead to improved long-term communication. The decision maker may also perceive value in having a convincing rationale for his choice, both in communicating his ideas to others and for his own peace of mind. There is evidence (Brown (1970), Longbottom and Wade (1973)) that this indirect value constitutes much of the value perceived by users of the technique. But while this indirect value is clearly important, it is also very difficult to measure since improvements in clarity of thought and communication are hard to evaluate even if they can be shown to have occurred. In this paper we restrict attention to the direct value of analysis as defined above.

1.3 Approaches to Assessing the Value of Analysis

1.3.1 Approach I: separate measurement of U_I and U_O - There are several methods we can suggest for measuring the difference between U_I and U_O, but perhaps the most obvious would be to find an assessment method for U_I and U_O separately

and to subtract the one from the other to calculate the expected value of analysis.

It is possible to make the separate assessments of U_I and U_O directly, that is, to produce numbers intuitively which seemed to represent the decision maker's beliefs. But this assessment might well be little more than an intuitive assessment that analysis is, or is not, worthwhile in a particular instance. Moreover, the decision maker may well have little experience on which to base such an assessment and so may have very little confidence in the quality of the numbers produced.

A more promising approach, at least for the novice, may be to decompose the gambles of proceeding to a decision with and without analysis into their component events. For example, suppose that the decision is to select one of n alternative gambles. As a consequence of the analysis the decision maker will end up with the expected utilities of each of these gambles; let us call these $\{x_i\}$, $i=1 \ldots n$. Prior to carrying out the analysis, however, he is uncertain what these numbers will be. Because of his predisposition to the tenets of decision analysis, however, he is able to express his uncertainty as a probability distribution over the $\{x_i\}$; let $f(\underline{x})$ be its density. If, now, he is convinced that he will take that action which is shown, by the analysis, to have highest expected utility, his expected utility, prior to analysis, is

$$U_I = \int \max_i x_i \cdot f(\underline{x})\, d\underline{x}. \qquad \text{(a)}$$

This states that his expected utility of carrying out the projected analysis and then acting on the recommendations of that analysis prior to analysis is the expectation, with respect to his uncertainty about the outcome of the analysis, of the biggest expected utility that will be produced by the analysis. Similarly, we can derive an expression for U_O. If a decision maker takes the trouble to assess $f(x)$, he can easily deduce (at least in theory, if not in practice) his current expectations of what x_i will be; and thus if he proceeds with no analysis he will choose that action with the biggest current expectation. Thus,

$$U_O = \max_j \int x_j\, f(\underline{x})\, d\underline{x}. \qquad \text{(b)}$$

We have represented this situation by the decision tree in Figure 1-1.

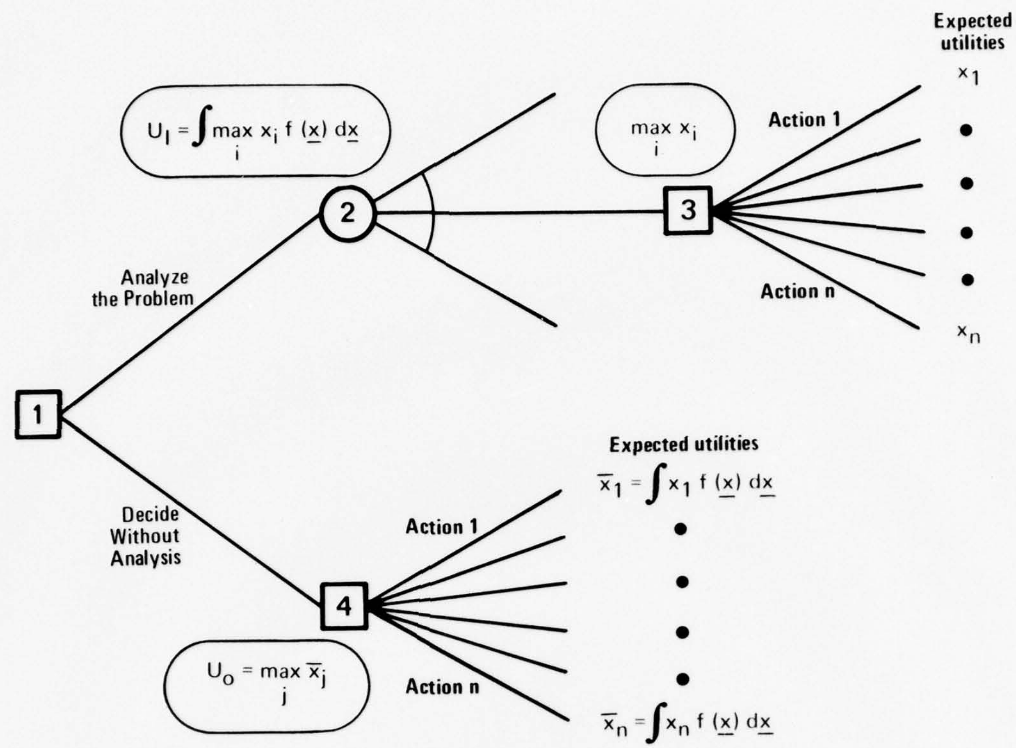

Notes:
(i) The expressions ringed at the nodes of this tree are the expected utilities there.
(ii) Node 2 represents the analysis which produces $\underline{x}$, the vector of expected utilities associated with the actions. $f(\underline{x})$ is the decision maker's density function for $\underline{x}$.

Figure 1-1

DECISION TREE FOR THE VALUATION OF ANALYSIS

The decision maker's assessment of U_I and U_O can now be derived from his assessment of $f(\underline{x})$. Of course, if he has any doubt that he will follow the indications of the analysis, these will not be good models of his beliefs, and he will want to use different expressions to calculate U_I and U_O; for example, in a previous report in this area (Watson & Brown (1975a)), we suggested replacing (b) by an expression involving separately assessed probabilities that any particular gamble would be taken in the absence of analysis. But let us suppose for the sake of exposition that (a) and (b) are acceptable models of the problem.

Even if this is the case, there are other difficulties with these expressions which arise from the notion of perfect and imperfect analysis. Since x_i is the expected utility produced by the projected imperfect analysis of gamble i, it may well not be the same number as would be produced by a perfect analysis.† We must distinguish between these quantities; define x_i^I and x_i^P to be the expected utilities attained from the projected imperfect analysis and from a perfect analysis, respectively; then we can imagine that our decision maker can express his uncertainty about the values of $\underline{x}^I$ and $\underline{x}^P$ prior to any analysis. Let $f(\underline{x}^I,\underline{x}^P)$ be his joint probability density function over $\underline{x}^I$ and $\underline{x}^P$. Then we should replace (a) by

$$U_I = \int x^P_{i*} \cdot f(\underline{x}^I \underline{x}^P)\, d\underline{x}^I\, d\underline{x}^P, \qquad \text{(c)}$$

where i* is the integer for which x_i^I is maximized. This formula implies that an imperfect analysis produces a set $\{x_i^I\}$ of expected utilities. The i*th option is selected, since it appears, according to the imperfect analysis, to have the largest expected utility. But its "real" expected utility (i.e. according to a perfect analysis) is x^P_{i*}. Similarly we should replace (b) by

$$U_O = \max_j \int x_j^P\, f(\underline{x}^I \underline{x}^P)\, d\underline{x}^I\, d\underline{x}^P. \qquad \text{(d)}$$

These expressions are a little alarming, given that our goal is to obtain a reasonably straightforward procedure for evaluating analysis prior to its execution. It would be tempting to ignore the distinction between perfect and imperfect analysis and merely use expressions (a) and (b) without further enquiry into whether the expected utilities produced are "correct" or not; but there is no general justification for this procedure. Without the benchmark of

†We are not suggesting here that perfect analysis is practically possible. We are merely using the concept as a benchmark in evaluating any practical analysis which almost surely could be improved.

perfect analysis, what is meant by asserting that the preference ordering produced by one imperfect analysis is "better" than a different ordering produced by a different imperfect analysis?

Fortunately, there is an extra condition on the decision maker's assessments which does allow us to simplify (c) and (d) to (a) and (b). Let us suppose that the density function $f(\underline{x}^I, \underline{x}^P)$ is such that

$$E(\underline{x}^P | \underline{x}^I) = x^I.$$

This formula seems very reasonable; for if, after an imperfect analysis, the expected value of the expected utility of an action that would be obtained by a perfect analysis differed from the expected utility actually obtained, the imperfect analysis could be instantly improved by changing u^I so that equality did hold here. Assuming that condition (e) holds, we can write

$$U_I = \sum_i \int_{A_i} \int_{\underline{x}_P} x_i^P \, f(\underline{x}^I, \underline{x}^P) \, d\underline{x}^P \, d\underline{x}^I,$$

where A_i is the set of points in $\underline{x}^I$-space such that $x_i^I = \max_j x_j^I$. Hence, taking the expectation with respect to $\underline{x}^P$, and writing $f_I(\underline{x}^I)$ for the marginal distribution on $\underline{x}^I$, we get:

$$\begin{aligned} U_I &= \sum_i \int_{A_i} x_i^I \, f_I(\underline{x}^I) \, d\underline{x}^I \\ &= \max_i x_i^I \, f_I(\underline{x}^I) \, d\underline{x}^I. \end{aligned}$$

This differs from (a) only in the use of the index I; but since in (a), all the quantities relate to imperfect analysis anyway, the two expressions are identical.

Similarly we can write (d) as

$$\begin{aligned} U_O &= \max_j \int E(\underline{x}_j^P | \underline{x}^I) \, f_I(\underline{x}^I) \, d\underline{x}^I \\ &= \max_j \int x_j^I \, f_I(\underline{x}^I) \, d\underline{x}^I. \end{aligned}$$

We thus recover equation (b).

We are now in a position to substantiate the assertion made in the previous section that the expected value of analysis can be negative. Of course, if condition (e) holds, (a), this assertion can never be the case since it is easy to see that U_I, as given by (a), cannot be less than U_O as given by (b). On the other hand, suppose that our decision maker is sufficiently incoherent that (e) does not hold. In particular, as an extreme example, suppose that he assesses $\underline{x}^I$ and $\underline{x}^P$ to be independent random variables. In this case, of course, the imperfect analysis is very bad, since its results give no indications about what would happen if a perfect analysis were to be performed. In this case, we can write (c) as

$$U_I = \sum_i \int_{A_i} \bar{x}_i^P f_I(\underline{x}^I)\, d\underline{x}^I$$

by using the same notation above, and where $\bar{x}_i^P = E(x_i^P)$. The formula

$P_i = \Pr\,[x_i^I > x_j^I$, for all $j \neq i]$ reduces to

$$U_I = \sum_i p_i \bar{x}_i^P.$$

Similarly, (d) may be written

$$U_O = \max_j \bar{x}_j^P.$$

Clearly $U_O > U_I$, and equality is only achieved if p_i is non-zero only for indices for which $\bar{x}_i^P$ is maximized.

Although this example is extreme, it illustrates our main point: if we have only a limited belief that an imperfect analysis will produce an answer close to the "correct" answer (that is, the answer produced by perfect analysis), then the expected value of that analysis can be negative. We suspect that one of the reasons why formal quantitative methods are not very popular with managers is for this very reason; they suspect that a projected analysis will be so badly done that it has a negative expected value.

1.3.2 Approach II: reduction in the expected cost of irrationality - A second, and very different, assessment method for the value of analysis, on the other hand, uses U_P, the expected utility from carrying out a perfect analysis and acting on its recommendations. One of the ways of assessing the value of partial information in decision analysis is to assess separately the value of perfect information and the expected opportunity loss because of unreliable

information. Analogously, we can calculate the value of a projected imperfect analysis from measurements of the value of perfect analysis and of what we might term the expected irrationality cost. (See Brown et al. (1974), p. 359 for an earlier discussion of this notion.) Remembering that we have specified the case when the decision maker's utility function for out-of-pocket expenses is linear, he is concerned to measure

$$U_I - U_O = (U_P - U_O) - (U_P - U_I). \qquad \text{(f)}$$

Our experience suggests that some decision makers might find the separate assessments of $U_P - U_O$ and $U_P - U_I$ more accessible than direct assessments of U_I and U_O. We have not yet investigated suitable ways to decompose such assessments into simpler judgements, but no doubt ways exist.

1.4 Implementation of Approach 1

In this section we further investigate ways of using (a) and (b) to measure the value of analysis. Note first that we have not mentioned the scale for the utilities involved. As is well known, a von Neumann-Morgenstern utility function is unique only up to a positive linear transformation; thus we can fix two points on it arbitrarily. In this description we shall take 0 to be the utility of the worst conceivable outcome to the decision, and 1 as the best; these definitions are purely for convenience and have no relation to the rest of our argument, but they imply that the expected utilities produced by the analysis will all be in the range [0,1].

To illustrate these formulae, let us suppose that we are confronted by a decision involving four alternatives. In contemplating a decision analysis, we realize that it would result in expected utilities x_1, x_2, x_3, x_4 associated with each of four alternatives. Currently we do not know what these expected utilities would be, but, being practised as subjective probability appraisers, we are able to express our uncertainties about these numbers by a multi-variate density function $f(\underline{x})$.

Let us for a moment put on one side the question how we might assess this density function in practice. It is, in fact, a formidable technical problem even if the objects of assessment are much more accessible, in the psychological sense, than those we have to consider here. But to illustrate the workings of our procedure, let us suppose that after a proper assessment of our beliefs, we conclude that they are represented by the density function,

$$f(x) = \prod_{i=1}^{4} 6\, x_i\, (1-x_i). \qquad \text{(g)}$$

Note that this density function has two properties that are unlikely to be in true practice; namely, it implies that the marginal distributions on each x_i are the same, and that the $\{x_i\}$ are statistically independent.

We can now insert (g) into (a) and (b) to evaluate U_I and U_O. Rather than directly evaluating the integral in (a), we recall that for a set of n independent random variables with a common cumulative distribution function F(x), the maximum of the set has a cumulative distribution $[F(x)]^n$. Thus, we can write (a) as

$$U_I = \int_0^1 z \cdot 4[3z^2-2z^3]^3 \cdot 6z(1-z)dz$$
$$= .733.$$

Clearly an application of (4.1) to (3.2) gives

$$U_O = .5.$$

To highlight the meaning of these calculations, suppose further that the only attribute of value in the decision is the monetary outcome, that the spread of possible outcomes is £1m, so that u(£1m) = 1 and u(£0) = 0, and that utility is linear in money. These assumptions imply that our current expected utility to be gained by going ahead with the analysis and acting on its recommendation, is equivalent to a monetary sum of £733,000, while the option of going ahead without analysis is only equivalent to £500,000. Thus, in this case, the value of analysis is £233,000, and it is irrational to spend more than this amount in analyzing the decision.

This direct approach has several severe drawbacks; in particular there are enormous difficulties in making assessments such as (g), and there are also great computational complexities in calculating U_I and U_O given that assessment. (For example, there is no simple analytic form for (a), if $f(\underline{x})$ is a distribution as well-used as a multivariate normal.) We now give a graphical representation of these ideas which gives insights into approximate methods for assessment and calculation.

Let us take another example; suppose that there are just two possible actions which we can take and that as a consequence of the proposed analysis, expected utilities x_1 and x_2 will be produced for these actions. Then let us suppose that we are able to produce marginal distributions on what each of these numbers will be; we have plotted cumulative probability distributions for these in Figure.1-2.

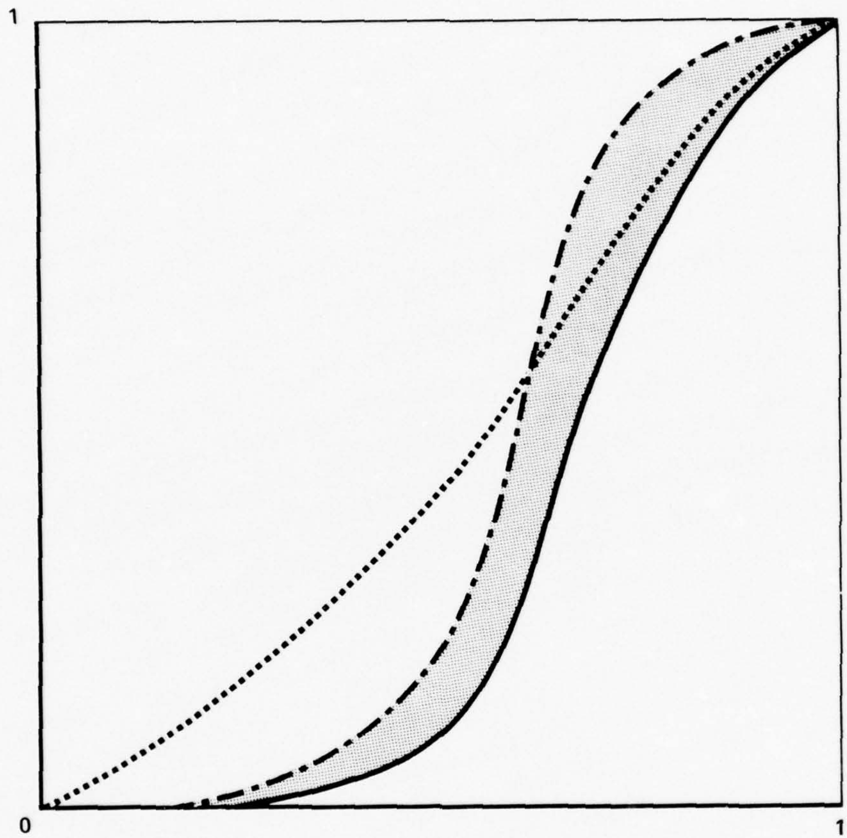

Cumulative distribution for x_1
Cumulative distribution for x_2
Cumulative distribution for max (x_1, x_2)

Figure 1-2
A GRAPHICAL METHOD FOR EVALUATING DECISION ANALYSIS

Now for a random variable defined on [0,1], the expected value is the area above its cumulative probability curve within the unit square. Thus U_O, measured according to (b), is the largest of the areas A_1, A_2, where these are, respectively, the areas above the cumulative curves for x_1 and x_2. In our example, it is clear that A_1 is the larger, so that U_O is equal to the area A_1. Suppose now, as in our previous example, that our beliefs about x_1 are independent, in the statistical sense, of those about x_2. Then the cumulative for the maximum of these two is just the product of the independent cumulatives, as indicated in the figure. In this case, U_I given by (a) is just the area above this cumulative curve. The value of analysis, U_I-U_O, is just the size of the area between the cumulative curve for X_1 and the cumulative curve for max (x_1,x_2).

Accordingly, in this case the value of analysis can be measured directly without further ado. Moreover, the extension to deal with more than two alternatives is straightforward and immediate. The significance of this method is that by representing the problem in this way, we can easily assess the change in our estimation of the value of analysis that would result from modification of any of our inputs. To be specific, we give some examples.

i. If, for example, we revised our beliefs about what x_1 would be since we were less sure about this number than the cumulative curve in Figure 1-2 would indicate, then this curve would flatten out, affecting the cumulative for max (x_1,x_2) as well. The area between them, which represents the value of analysis, would also change, but it would be straightforward to see by how much and in which direction.

ii. Again, if we felt that x_1 and x_2 were positively correlated, the cumulative for max (x_1,x_2) would lie above the curve drawn to represent it in Figure 1-2. The extent of the correlation would determine by how much the curve would rise, but it would be bounded above by the smaller of the other two cumulative curves, since

$$\Pr\,[\max\,(x_1,x_2) \le x] \le \min \left[\Pr\,[x_1 \le x],\ \Pr\,[x_2\ x]\right].$$

iii. It may be that we are unhappy about representing U_O by (b). Perhaps in the absence of analysis, we will decide by tossing a coin. In that case, a more appropriate expression for U_O would be

$$U_O = 1/2 \int x_1\ f(\underline{x})\ d\underline{x} + 1/2 \int x_2\ f(\underline{x})\ d\underline{x}.$$

This formula gives the area above the curve formed by taking the average of the marginal cumulative curves for x_1 and x_2. The area between this curve and the cumulative for $\max(x_1,x_2)$ is, once again, the expected value of analysis.

This graphical method provides considerable insight into the expected value of a projected analysis, and we have used it to give a very rough estimate of the value of analysis after only a very limited amount of reflection. Further research is needed, possibly using interactive computer graphics, to develop these ideas.

1.5 Towards Practical Implementation

There are, clearly, several important problems to be faced before we can have a great deal of confidence that any of the approaches outlined in section 1.3 constitute adequate measures for a decision maker who is wondering whether to use analysis or not. Take, for example, the algorithm explored in section 1.4. We have assumed that the marginal distribution on each x_i can be assessed with confidence. But many people have a great deal of difficulty in assessing a probability distribution on a number as definite as, say, the population of London. When the number is merely that which is to be produced by an ill-defined projected analysis, the difficulties are compounded. Moreover, this approach makes the assumption that the decision problem is sufficiently clear-cut that the number of alternative actions open to the decision maker is well-known. Practiced decision analysts, on the other hand, report that a major part of many studies is the specification of the set of alternative courses of action. This finding may suggest that the other approach to valuing decision analysis, using expected irrationality loss, may be worth pursuing since, at least in the state of development outlined at the end of section 1.3, with this approach we do not need to specify how many alternative actions are contemplated. It does have the disadvantage, however, because the judgments required are at such an intuitive level, that they are, in some sense, rather distant from the details of the decision problem in hand.

Since we have been suggesting practical ways to assess the value of analysis, prior to its use, it is important to develop some experience in applying these approaches in practice. In a previous report (Watson and Brown (1975b)), we analyzed three cases of the application of decision analysis to assess the value of that analysis after it had been performed. Clearly, analyses of this kind should provide a body of experience on which to base judgments about the value of future analyses, but they will be more useful if a prior estimate of value has also been made. We are currently engaged on a program of research to make such

estimates in a number of cases, and we hope that the insights we gain will enable us to present a further study discussing which, if any, of the approaches outlined in section 1.3 are most promising as tools for the practical valuation of analysis.

The critical reader might accept that the framework for evaluation we have presented here is reasonable, but he may doubt whether it could ever constitute a practical device for guiding decisions about the allocation of resources to decision analysis, given that the subjective judgment required seem so inaccessible. On the contrary, while not asserting that our method can evaluate analysis even to one significant figure, our limited experience with these methods suggests that considerably more confident judgments can be made about the amount of money it is worth spending on the analysis of a particular problem by using our framework than can be made without it.

REFERENCES

Brown, R.V. (1970), "Do Managers Find Decision Theory Useful?" Harvard Business Review, 48 (3), 78-79.

Brown, R.V., Kahr, A.S., and Peterson, C. (1974), "Decision Analysis for the Manager," Holt, Rinehart & Winston, New York.

Brown, R.V. and Lindley D.V. (1976), "Rationality and the Resolution of Incoherence," unpublished working paper, University College, London.

Hampton, J.M., Moore, P.G., and Thomas H. (1973), "Subjective Probability and its Measurement," J.R. Statist. Soc. A 136, 21-42.

Hull, J., Moore, P.G., and Thomas, H. (1973), "Utility and its Measurement," J.R. Statist. Soc. A 136 226-247.

Lindley, D.V. (1971), "Making Decisions," Wiley-Interscience, New York.

Longbottom, D., and Wade, G. (1973), "An Investigation into the Application of Decision Analysis in United Kingdom Companies," Omega, 1, 207-215.

Matheson, J.E. (1968), "The Economic Value of Analysis and Computation," IEEE Journal on Systems, Science and Cybernetics, SSC-4, 325-332.

Raiffa, H. (1968), "Decision Analysis," Addison-Wesley, Reading, Mass.

Rice, T.R. (1974), "The Economics of Decision Making," Doctoral Dissertation, Stanford University, Department of Engineering-Economic Systems.

Watson, S.R., and Brown, R.V. (1975a), "Issues in the Value of Decision Analysis," Technical Report 75-9, Decisions & Designs Inc., McLean, Virginia.

Watson, S.R., and Brown, R.V. (1975b), "Case Studies in the Value of Decision Analysis," Technical Report 75-10, Decisions and Designs, Inc., McLean, Virginia.

2.0 A CASE STUDY IN THE VALUATION OF DECISION ANALYSIS

By S. R. Watson

In recent papers (1, 2, 3) we have observed that the problem of how much effort to spend in analyzing a decision is itself a problem of decision and should therefore be amenable to formal decision analysis. In (1) and (3), we suggested ways to discover the expected value of a proposed analysis prior to undertaking it; in this report, we discuss the application of these methods to a particular case. We shall assume a familiarity with the ideas discussed in these papers, and we shall make extensive references to (3). Many methodological problems arose in our attempts to evaluate the analysis; rather than first describe the case and then discuss these problems, we shall give a continuous narrative interrupted where necessary by italicized passages discussing each methodological problem as it arose.

In the second half of 1976, one of the authors (R. V. Brown) held a fellowship funded by the Social Science Research Council, of the U.K., to encourage the use of decision analysis in the U.K. In this role, the Technical Director of Universal Foodstuffs (U.F.), Mr. R. Williams asked Dr. Brown to apply the methods of decision analysis to a particular problem then facing U.F. This problem was whether or not to increase or decrease manufacturing and marketing effort in a relatively new, but proven, food additive of wide applicability in the food industry. U.F. was currently manufacturing these substances in rather small quantities. At this stage, the authors decided to work jointly on the analysis of this decision together with some of the staff of U.F. Additives Division. Before we describe the development of the analysis further, we make the following observations about assessing the value of this analysis.

Note 1 - The motivation of the theoretical development in (3) was the need to give a decision maker some benchmark for deciding whether to spend a given amount of money on the analysis of a particular decision. In this case, the motivation was absent since the authors were advising in this analysis as a part of a research program and they were not intending to receive a consultant's fee from U.F. Our attempts to measure the value of this analysis are therefore necessarily made in a slightly different context than that we had in mind in developing the formulation. Moreover, at no stage in the consideration of the value of this analysis did we discuss with Mr. Williams his perceptions of what that value was. Nonetheless, an acceptance of the framework put forward in section 1.3 of (3) implies that Mr. Williams must have possessed the perceptions which would imply what

value this analysis had; accordingly, it makes sense to discuss this quantity. As a surrogate for the perceptions of the decision maker, we used our own perceptions, a fact which must be constantly borne in mind in what follows.

Note 2 - We made a distinction in section 1.2 of (3) between the direct value of analysis and the indirect value; the former was the extent to which analysis produced a better decision, in the particular instance of application, at least in expectation, and the latter was all other aspects of value derived from using analysis, such as the permanent improvement of communication within the organization. In this case we do not know which of these aspects of value Mr. Williams hoped to gain by using the technique; we shall try to consider only a measure for the direct value of the analysis by following a restriction similar to that in (3).

Note 3 - The structure for valuing analysis presupposes, as does decision theory itself, that there is a single individual making decisions based upon his values and perceptions of uncertainty. On the other hand, it is a commonplace that decisions made in organizations often emerge gradually and collectively; a consensus grows for a particular course of action without any obvious moment when the decision is taken and with no obvious decision maker. This case appeared to be of this kind, in that discussion of the issue had been underway in Additives Division for some time. But it is clear that Mr. Williams needed to decide what he thought the Division should do, and it may be that one of the reasons he welcomed the approach of decision analysis was to clarify his understanding of the issues. We can view our present exercise as an estimation of Mr. Williams' view of the value of this exercise to U.F.

On Tuesday, 6 July 1976, we had a preliminary discussion of the problem with Mr. N. Askwith of Additives Division and some of the staff of U.F. Central Management Services. Mr. Askwith gave a detailed description of the varied markets for the new additive, but in his view, there was very little scope for choice. He felt that a major development in the business was clearly called for and that the only open question was whether it should be by acquisition of other companies or by internal expansion. He also felt that even this question was virtually settled since there were very few other companies that could possibly be acquired by U.F.

Note 4 - At this stage we felt that this particular decision did not need decision analysis, since there was very little doubt about what to do. Heuristically, then, we felt the value of analysis to be very low. But this conclusion would have been valid only if we were convinced that the perception of the problem we then had was the one held by Mr. Williams, and we felt this was unlikely to be the case.

Note 5 - We did not attempt a formal analysis at this stage since we were not at all sure what the problem was. Now the argument which suggests that there is, for each decision maker contemplating analysis of a decision, a unique certainty equivalent for the value of that analysis, remains valid no matter how little the decision maker knows about the decision he has to make. Thus, we should have been able to use our framework to evaluate the analysis at this stage. We explain the fact that we felt such an exercise would be meaningless by pointing out that the less known about a decision problem, the greater the variance of measurement error in any assessment of the crucial quantities U_I and U_O. With almost no knowledge, the fact that U_I and U_O will be measured very imperfectly removes any confidence that their difference is at all close to the actual value of analysis at this stage.

We had further meetings with some staff of Additives Division on Tuesday, 10 August and Friday, 13 August. At those meetings it became clear that Mr. Williams did not perceive a major expansion in the new additive as the only possible decision. His brief to Additives Division was to decide at what level to continue manufacturing these chemicals. The Chairman of Additives Division endorsed the view that, although the space of decision options was continuous and rich, it made sense to analyze the consequences of only four different strategies, choosing these to span, in some sense, the space of strategies. These were:

1. Complete withdrawal from the new additive over a three-year period,

2. Continuation of sales along the historical-trend line,

3. The current proposals for development of the market, and

4. An all-out effort to gain the largest possible business in all seventeen end-use areas.

It was decided to build a simple cash-flow model by taking as inputs estimates for sales, prices, unit variable costs in each product and geographical area for the years 1976, 1981, and 1984, and then calculating trading profit for those years. This model would be used in a Monte Carlo simulation to derive probability distributions on these quantities. In addition, distributions on the capital requirements would be provided for each strategy.

Note 6 - The ideas of section 1.2 of (3) do not presume any particular structure for the decision problem, and it is quite general to assert that the value of an analysis is U_I-U_O.

The measurement of these numbers, however, requires some sort of decomposition, and the approach to decomposition described in some detail in section 1.3.1 of (3) makes the assumption that the decision problem has a finite list of alternatives which are known at the time that the analysis is valued. Not only will this method be inapplicable when it is far from clear what the options are; it also presents some difficulties when, as in this case, it is clear what the options are, but they constitute a continuum. If we consider a finite set of options, then Approach I will determine the value of analysis given an initial decision to restrict action in any case just to members of this set. As long as the option actually chosen and the option that would have been chosen in the absence of analysis are quite close (in the sense of expected utility) to some particular members of the finite set of options considered, then Approach I will give a good approximation to $U_I - U_O$. Thus, an extra inaccuracy in measuring the value of analysis arises from the simplification to a discrete set of alternatives.

At this stage we applied the formulae of Approach I of (3) to estimate the value of this analysis. Note first that the analysis as proposed falls short of the classical structure for such analyses in failing to combine the different attributes of value into one single utility measure. However, we supposed that net present worth would suffice for such, and we further supposed that the utility of U.F. was linear in net present worth. We were then faced with the problem of producing our own joint probability distribution on the numbers that would be produced by this analysis if the extra step of producing a single expected utility were to be carried out.

In this we were severely hampered by having very little idea of the production costs or capital investment necessary to carry out any of the strategies. The only data to hand at that stage were sales targets for the years in question. Current feeling was that if strategy 3 were continued, total worldwide sales of the additive £3.0M in 1976, £8.3M in 1981, and £12.2M in 1984. Since major U.K. companies have Net Trading Profits of the order of 1%-10% of sales turnover, we took 4% as a reasonable figure and found NTP of £120K, £332K, and £488K for the years 1976, 1981, and 1984.

To get some idea of what sort of figure might be produced by the financial model, we projected a cash flow that might reasonably follow if strategy 3 were pursued. The present value discounted at 15% of net trading profit streams over the next 20 years and whose trading profit values in the years 1976, 1981, and 1984 are the figures given above is about £2M. Supposing capital investment to be about ten times annual trading profit, we would expect investment of about £1M, giving a net present worth of £1M. Using arguments

of this kind, we derived our cumulative distribution for the net present worth figures that would be produced by the analysis for each of the four strategies, and these are plotted in Figure 2-1. We also felt it to be a reasonable approximation that our assessment of one of these numbers did not depend on the value of others, that is, that these variables are statistically independent. Since we felt very sure at that stage that strategy 3 would have been pursued if there had been no analysis, the value of analysis is the area between the curve for strategy 3 and the cumulative curve for the maximum of these random variables, which we have labelled (5) in Figure 2-1. The value of analysis given by the hatched area is approximately ₤1.7M.

Note 7 - We now asked how accurate this measure might be. If the cumulative curves in Figure 2-1 were moved a little, they would still be acceptable representations of our degree of belief in what the results of the analysis would be. This difficulty occurs in much more straightforward assessment problems; for example, we often find it difficult to distinguish beyond the first decimal place (or even as accurately as that) when assessing the probability of a single uncertain proposition. The conventional way to overcome that difficulty is to decompose the assessment in several ways until the assessment has been sufficiently refined. If we were to adopt such a procedure here, then the effort involved would be enormous and the whole aim of the exercise would be vitiated, for what we are seeking is simple guidelines for how much to spend on a proposed analysis. Nevertheless, we can deduce what the order of magnitude of the value of the analysis is. For it to be as little as ₤200,000, the cumulatives in Figure 2-1 must be much steeper; that is, we must be much surer about the actual outcomes of the analysis, or our beliefs must be highly positively correlated (for this reduces the value of analysis), or we reduce our prior belief that strategy 3 will be adopted if we do no analysis. (If we adopt the decomposition of U_O given by equation (b) of (3) and keep the cumulative curves of Figure 2-1, then strategy 4 will be adopted if no further analysis is carried out, and the value of analysis reduces to about ₤1.2M.) Alternatively, negative correlations or a greater uncertainty in the values to be produced by analysis could result in larger figures for the value of analysis. To achieve a figure as high as ₤5M, however, the hatched area in Figure 2-1 would have to cover half the area between 0 and 1 on the abscissa and 0 and ₤10M on the ordinate; it is very difficult to see how the area could be increased to that extent, while maintaining marginal cumulative probability curves which represent satisfactorily our degrees of belief in how the analysis would unfold. Thus, we conclude that the value of the analysis at this stage was somewhere in the range of ₤200,000 to ₤3M.

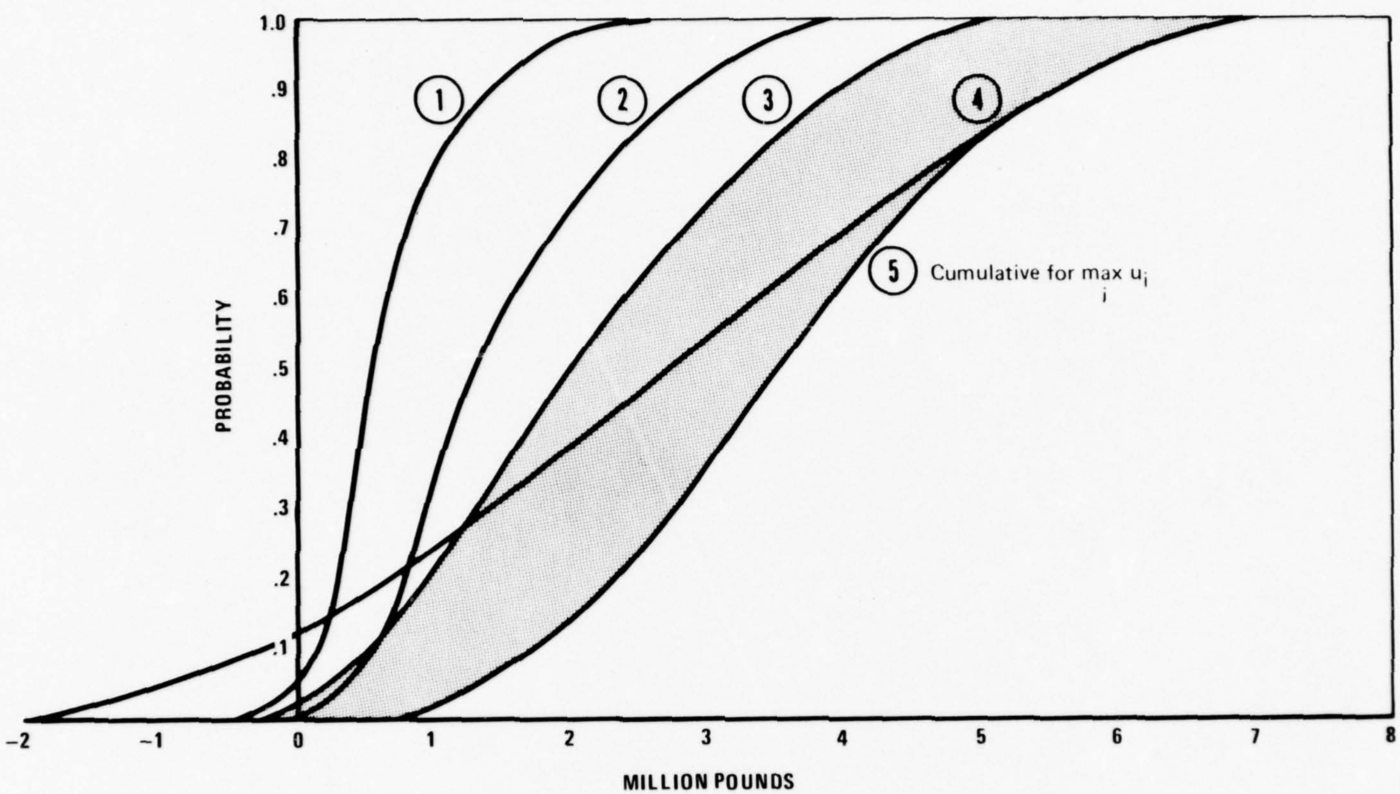

Figure 2-1

CUMULATIVE DISTRIBUTIONS FOR $\underline{u}$ IN UF CASE – FIRST ASSESSMENT

In the second half of August 1976, the technical staff at U.F. Additives Division continued the analysis, and on 6 September we visited them again to discuss developments. It was not possible to attempt a second evaluation of the analysis now that it was half complete. We had greater information about costs and likely revenues, and a better appreciation of what the company might have done if there had been no attempt to use decision analysis.

Note 8 - In this study we are contrasting the action taken by a company consequent to analysis, with the action taken in the absence of such analysis. We are not suggesting by this distinction, however, that any company, U.F. in particular, would actually take an important investment decision without analysis of any kind. Clearly the value derived from using the formal technique of decision analysis is likely to be much greater if the alternative decision-making procedure is random choice, shall we say, than if it is some other kind of analytical technique. All we are searching for is the additional value obtained by using decision analysis in excess of that obtained by using whatever other technique would have been used.

At this stage, we decided to evaluate the analysis by using both Approach I of (3) and also by using the method based on expected irrationality loss (Approach II), outlined in section 1.3.2 of (3). We first applied Approach I. Once again, we assessed our joint distribution for the figures of net present worth to be produced by the analysis, which we took to be linearly related to utility. On this occasion, we felt that these numbers were not statistically independent; on the other hand, we felt that they would be well represented by a 4-variate normal distribution. The marginals of this distribution for each of the four strategies were therefore also normal, and their means and standard deviations were assessed to be, respectively, .6, .12; 2.6, .8; 12.5, 5.8; and 15.0, 21.7. We have plotted these curves in Figure 2-2. To assess the correlation coefficients presented great problems of elicitation; however, it was clear that the exact values of these were not very important, and we felt it reasonable to make them all equal. To calculate the correlation coefficient from two assessments of the expected value of one variate conditional on the others, we used the fact that for a multi-variate normal distribution, the expectation of one variate conditional on the others is a linear function of the others. A value of .1 for the correlation coefficient emerged. The explicit calculation of the cumulative distribution of max u_i if $\underline{u}$ has a multi-variate normal distribution is highly complex either analytically or numerically, but there exist approximations to carry out this calculation reasonably straightforwardly. We used these approximations to calculate the cumulative for max u_i in this case, and the result is plotted in Figure 2-2. At this stage, we revised

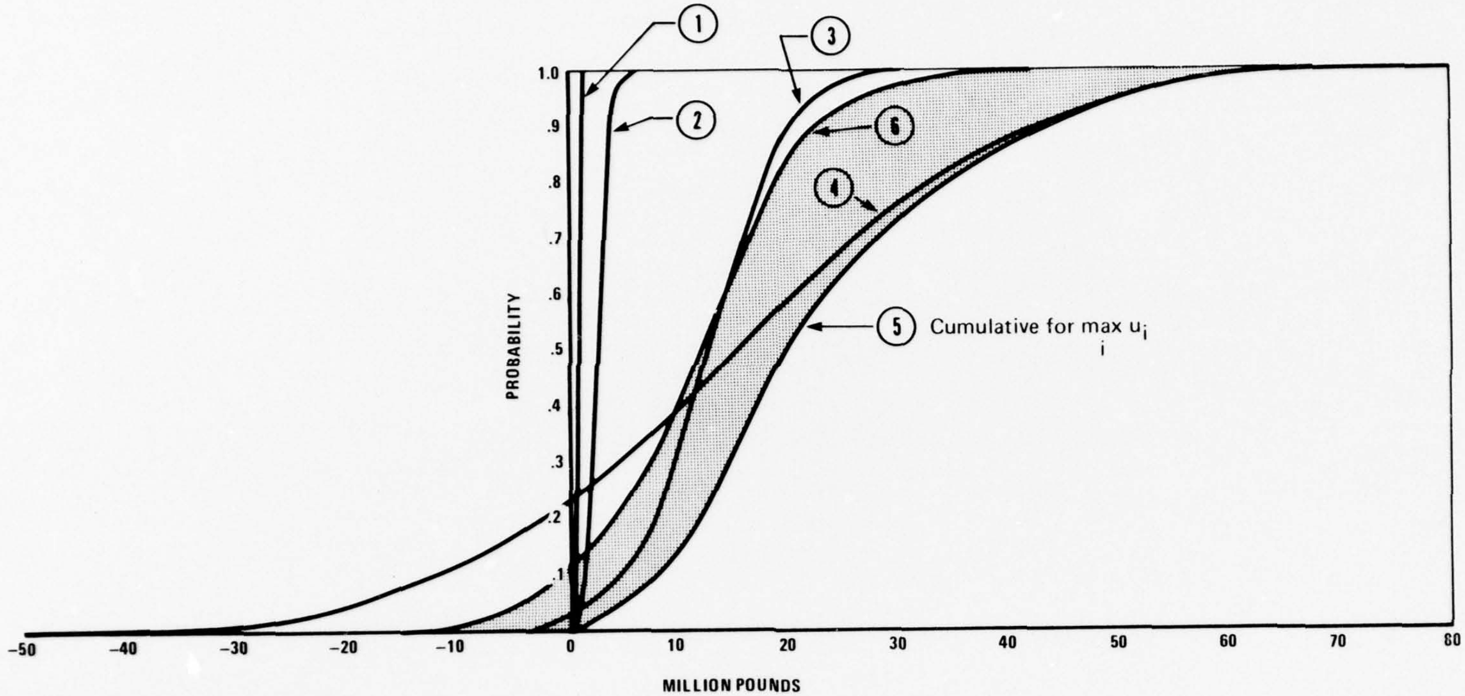

Figure 2-2

CUMULATIVE DISTRIBUTIONS FOR u IN THE UF CASE – SECOND ASSESSMENT

our ideas about the decision that would be taken in the absence of analysis. We assessed the probabilities that strategies 1-4 would be taken in the absence of analysis as .01, .07, .75, and .17, respectively. Thus, using the decomposition for U_O given in (1) (rather than that in (3)), we drew in Figure 2-2 the cumulative for $.01u_1 + .07u_2 + .75u_3 + .17u_4$; this is the curve marked (6) in that figure. The value of analysis is the area between this curve and the cumulative for max u_i, marked (5) in Figure 2-2. The area corresponds to approximately ₤10M. This judgment is considerably in excess of the previous one, and its size results from the higher probability that strategy (3) will be taken without analysis than the probability that the analysis will show (4) to be a considerably better strategy.

We also used Approach II to evaluate the analysis. This involves a separate assessment of U_P-U_I, the expected cost of irrationality, and U_P-U_O, the expected loss of value because of imperfect analysis. Some brief contemplation suggested that the present value of net cash flow to U.F. might well vary between +₤80M and -₤50M. With these figures in mind, we felt that the expected cost of irrationality was unlikely to be more than ₤20M, and indeed might well be very small, particularly if the analysis were to show that the action originally contemplated was indeed the best. On this qualitative basis, we made a direct assessment of the expected cost of irrationality to be

$$U_P-U_O \quad = \quad ₤7M.$$

Given the limitations of the analysis that was in fact performed, which we felt to be quite severe, we assessed that about ₤5M could be expected to be lost through imperfect analysis. Thus, with U_P-U_I = ₤5M, we deduce the value of analysis to be ₤2M. This amount is of the same order of magnitude as we derived by the Approach I; the obvious shortcomings in the assessment techniques used by both approaches prevent us from concluding at this stage that one approach has greater validity than the other.

Conclusions

Our experience in this case reassures us that our approach to the value of analysis gives a framework in which we can assess how much it is rational to pay for such an analysis; in this case, we can be confident that the value is at least hundreds of thousands, if not millions, of pounds. The elicitation techniques we have developed so far are not very accurate, but they are at the least a starting point for a more effective procedure. We can make qualitative observations about the factors that affect the value of a proposed analysis; naturally, these are in accord with our intuition. Value is large if

(i) the stakes involved are high,

(ii) there is a large probability that as a result of the analysis the action selected will change, and

(iii) there is initially great uncertainty about which is the best option to take.

In the Notes within this paper, we have mentioned some of the methodological difficulties we faced; we feel, however, that these do not detract from the fundamental ability of our procedures to provide a framework to assess how much effort it is worth expending in analyzing a particular problem.

REFERENCES

1. S. R. Watson and R. V. Brown (1975), "Issues in the Value of Decision Analysis," Technical Report No. 75-9, Decisions and Designs, Inc., McLean, Va.

2. S. R. Watson and R. V. Brown (1975), "Case Studies in the Value of Decision Analysis," Technical Report No. 75-10, Decisions and Designs, Inc., McLean, Va.

3. S. R. Watson and R. V. Brown, "The Valuation of Decision Analysis," paper 1 in this report.

3.0 NOTES ON RATIONALITY AND THE RECONCILIATION OF INCOHERENCE

by R. V. Brown and D. V. Lindley

3.1 The Problem

A government official was recently faced with the problem of making his probabilistic assessment of various components of US energy demand in order to make an energy policy decision. Each "target" quantity could be expressed, in a number of alternative ways, as a function of different components, for example, as the product of number of users times consumption per user or as the products of number of equipment units times consumption per equipment unit summed over equipment types. From surveys and other sources, he made probabilistic estimates for each component and, by using decision theory, found the implied probability distribution for the target quantity. However, the different decomposition functions yielded quite different probability distributions for the same target and, derivatively, quite different energy policy decisions. On what distribution should he base his decision, and how can he justify that choice? If one choice is as good as another, why should he not simply make a direct assessment of the target and disregard any of the more sophisticated approaches to uncertainty assessment which decision theory permits?

Is there only one logical way for a subject, S, say a decision maker, to process what he has in his head at any point in time? Is there only one "right" decision or inference for him on any issue in terms of his total field of perception? Established decision and probability theory tells him whether any particular set of his assessed probabilities, utilities, and act preferences are logically coherent. They will test whether the act he subjectively prefers is indeed the one with the highest expected utility according to a particular set of his probability and utility assessments. They will test whether his directly assessed probability for an event after learning new information is what Bayes' Theorem computes it to be from his prior probability and his "likelihood function" for that information. And so forth.

If S's assessments always did cohere, decision theory would have little practical value for him since his direct assessments could not be improved upon. Fortunately for decision analysts, psychologists have established that subjects are typically incoherent and have even derived measures of incoherence.

What should the demonstrably incoherent subject do if he wishes to be rational? Is there some unambiguous and compelling principle by which his incoherence can be resolved? Of course, any system of assessments can be made coherent, by arbitrarily adjusting their values. But is any one such reconciliation superior to another and on what grounds? Coherence by itself does not seem to provide an answer, and, without one, the very foundations of decision theory as a prescriptive tool are challenged.

3.2 An Approach Based on Assessment "Validity"

It is intuitively appealing to argue that some of S's "raw" assessments are, in some sense, more "valid" than others, and that this relative validity should somehow be taken into account when reconciling initial assessments. A satisfactory procedure appears to have four requirements:

1. A definition of the validity of initial raw assessments,

2. An explicit specification of alternative reconciliation methods,

3. A paradigm to derive reconciled assessment validity from raw assessment validity, and

4. A basis for choosing between reconciled systems in terms of validity.

A suggestive approach to defining assessment validity would be to define it as a measure of analytic stability. Raw assessments could be characterized by a probability distribution on "shift on further reflection" (not with further information). (More generally, this would be a joint validity distribution reflecting "shift" dependencies between assessments.) Variance might be a useful summary measure of validity.

An irksome problem here is how to define "shift on reflection" without taking for granted the optimal resolution of incoherence which the "shift on reflection" itslef is to be an instrument in discovering. Even if there is some degree of such circularity in definition, perhaps the "further reflection" can be specified as an uncertain procedure whose expected impact can nevertheless be assessed. Some allowance must also be made for S being incoherent in assessing his validity distribution.

Any particular reconciliation procedure for mapping raw assessments onto coherent assessment space might be interpreted as a function (for example, pooling different direct

and indirect assessments of the same target assessment as a weighted average with weights proportional to the precision of the component assessments).

Established probability calculus will determine the distribution of any functions of random variables. Presumably the validity distribution for a reconciled system can then be derived from the validity distribution of the raw assessment system.

The preferred reconciliation method is the one which maximizes some measure of total system validity. (This would not necessarily give unique rationality, but it would give maximum rationality.) No persuasive single scalar measure of validity is apparent to characterize a completely (or a common) reconciled system. However, the reconciled system can be optimized with respect to, say, the validity variance of any one specified target assessment. Therefore, the optimal system and, the optimal method of reconciling initial incoherence may depend on which the target assessment is. We may be able to say only that we know in principle how to resolve incoherences for the purpose of a single target assessment, but we may have to acknowledge that a different reconciliation may be appropriate if a different target assessment is involved. Unique rationality therefore escapes us.

A possible approach to system optimization is to treat the choice of a reconciled assessment system as a decision, much as one might choose an estimate of probable product demand on which to base a business stocking decision. However, the analogy appears brittle when probed. It is not clear how one would define the "loss structure" called for in this type of problem. There appears to be no constructive analogue to the true value by which the loss is to be defined, much less to a probability distribution on divergencies from that value.

3.3 A Bayesian Updating Approach

An alternative approach to the resolution of incoherence involves positing an investigator, N (distinct from the subject, S), who is to determine a unique, rational system of assessments for S. Unlike S, N is treated as perfectly coherent. He has a prior distribution on S's rational system of target assessments, T, and a likelihood function for T, given S's raw assessment system, R. Through Bayes' Theorem he can derive a posterior distribution on T. It is probably mathematically demonstrable that the variance of N's posterior on T gets smaller, possibly to the point of vanishing, as R is extended to include more and more of the subject's field of perception. (How fast, depends on the

diagnosticity of the likelihood function.) This much is investigator independent and confirms one's intuitive conviction that it pays to address a target assessment in as many different ways as possible (much as it pays a surveyor to take many different bearings on a location).

However, we are left with the problem of having a reconciliation procedure dependent on characteristics attributed to the investigator. Where do N's priors and likelihood functions come from? Since N is a hypothetical construct, they should not be idiosyncratic to N but should somehow be descriptive of S. Can we treat S as a coherent assessor for this purpose? Can Perhaps any incoherence here causing second-order fuzziness perhaps be disregarded?

If the likelihood function is informative enough, any "uninformative" prior, however defined, may lead to virtually indistinguishable results and so be acceptable. But according to what principle should N (or S) construct the likelihood function? It would appear to be closely related to the raw validity distribution called for by the first approach and appears subject to the same conceptual problems.

With either approach, S's incoherent views (on assessment validity and likelihood, respectively) could be taken into account by N by Bayesian updating, much as he would take into account any expert probability distribution in forming his own probability distribution. In this way, perhaps we can maintain conceptual distinction between S and N with either approach, while exploiting to the full S's perceptions.

3.4 Glossary

Subject (S) - the person we are looking at (at a given point in time unless otherwise specified).

Object - a real world entity such as event, act, relationship.

Assessment - a value (for probability, utility, choice) subject assigns to an object.

Elicitation - the direct assessment of a quantity in a decision system (e.g., probabilities, utilities.

Target (T) - an object or assessment the subject is primarily interested in.

Perception Field (F) - everything in S's head - totality of actual or potential perceptions, and availablity for elicitation.

Raw Assessment (R) - a number (e.g., probability) elicited straight from S's Field (i.e., unconstrained by coherence).

Target Function (TF) - algorithm (e.g., Bayes Theorem) deriving target from other assessments (e.g., prior, likelihoods).

Target Space - set of target assessments implied by some set of raw elicitations.

Assessment Structure - group of related target- and cross-functions, a model without specific assessments but with potential incoherence.

Assessment System - quantified structure (i.e., with specific assessments.

Reconciled System - system with coherent assessments (vs. raw system).

Assessment Validity - measure of the rationality (e.g., firmness) of an assessment (raw or derived).

Assessment Stability - a measure of validity (e.g., variance of validity distribution of shift in assessment on further analysis).

Optimal System/Target - "most rational" target assessment (and embedding system) (e.g., with highest "validity").

Decomposition - expressing an object variable as a function of other object variables (e.g., demand = customer x demand per customer).

Coherence - logical compatibility (e.g., according to probability calculus).

Unique Rationality - the concept that subject has a single, most coherent interpretation of his field of perception.

AFTERWORD

Although the basic development of valuation techniques is by no means complete, it is at a stage where it is realistic to attempt the development of operational procedures which are economical to apply in the course of applied decision analysis. One such procedure, involving a simple graphical display, has been formulated. A start has also been made to build up a case file of hindsight evaluations from DDI project experience.

Under proposed follow-on research, work on the mathematical foundations of the value of decision analysis will build upon the research efforts by Watson and Brown reported here and will focus on identifying and proving key relationships that are needed for the development of applied valuation methodology. In addition, work will be pursued which focuses on the direct development of operational valuation algorithms for field application.

Under the current contract, efforts have been initiated to develop operational algorithms for evaluating decision analysis projects in conjunction with actual decision analytic applications. This work must be extended and expanded substantially to reach fruition. While the algorithms for prior valuation of analysis presented in Paper 1 have some theoretical appeal, and show some promise in relatively straightforward business applications, they have yet to be tried on defense and governmental situations where the multi-attributed valuation of consequences is a good deal more complex. It is clearly important to attempt to do this. We expect difficulties to arise in using these valuation algorithms in practice, which will be uncovered by attempting to use them on real cases. These difficulties will be in elicitation, in that decision makers might well find the probability judgments necessary for the algorithms rather hard to make. Psychological insight and skill will be necessary to discover good ways to perform such elicitations. There may also be some need to modify the formulae, so that the elicitations required become more accessible, but in such a way that the valuations obtained remain roughly the same as they would be if the proper formulae were used.

In the long run, it should be possible to distill the observations about value made from the study of the analytical process in many different cases into a set of generalizations that would be of use in evaluating further analyses.

A major attraction of the initial case analyzed in Paper 2 is that the dimensions of value of the problem are

primarily monetary, thereby simplifying the task of assessing a monetary value for the decision analysis.

It is now proposed to address further cases which are more ambitious in that, as in most defense and government decisions, the dimensions of value are not predominantly monetary. Ongoing case studies being worked on at DDI will provide the illustrative contest for methodological development.

OFFICE OF NAVAL RESEARCH, CODE 455
TECHNICAL REPORTS DISTRIBUTION LIST

Director, Engineering Psychology
Programs, Code 455
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217 (5 cys)

Defense Documentation Center
Cameron Station
Alexandria, VA 22314 (12 cys)

Dr. Robert Young
Director, Cybernetics Technology Office
Advanced Research Projects Agency
1400 Wilson Blvd
Arlington, VA 22209

Col. Henry L. Taylor, USAF
OAD(E&LS) ODDR&E
Pentagon, Room 3D129
Washington, D.C. 20301

Capt. Roger Granum
Office of Assistant Secretary
of Defense (Intelligence)
Pentagon, Room 3C200
Washington, D.C. 20301

Office of Naval Research
International Programs
Code 102IP
800 North Quincy Street
Arlington, VA 22217 (6cys)

Director, Naval Analysis Programs
Code 431
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director, Operations Research
Programs, Code 434
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director of Statistics and
Probability Program, Code 436
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director, Information Systems
Program, Code 437
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Commanding Officer
ONR Branch Office
ATTN: Dr. J. Lester
495 Summer Street
Boston, MA 02210

Commanding Officer
ONR Branch Office
ATTN: Dr. Charles Davis
536 South Clark Street
Chicago, IL 60605

Commanding Officer
ONR Branch Office
ATTN: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Dr M. Bertin
Office of Naval Research
Scientific Liaison Group
American Embassy, Room A-407
APO San Francisco 96503

Director, Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D.C. 20357 (6 cys)

Mr. John Hill
Naval Research Laboratory
Code 5707.40
Washington, D.C. 20375

Office of the Chief of Naval
  Operations, OP987P10
Personnel Logistics Plans
Department of the Navy
Washington, D.C. 20350

CDR Paul Nelson
Naval Medical R&D Command
Code 44
Naval Medical Center
Bethesda, MD 20014

Director
Behavioral Sciences Department
Naval Medical Research Institute
Bethesda, MD 20014

Dr. George Moeller
Human Factors Engineering Branch
Submarine Medical Research Laboratory
Naval Submarine Base
Groton, CT 06340

Bureau of Naval Personnel
Special Assistant for Research
  Liaison
PERS-OR
Washington, D.C. 20370

Navy Personnel Research and
  Development Center
Management Support Department
Code 210
San Diego, CA 92152

Dr. Fred Muckler
Navy Personnel Research and
  Development Center
Manned Systems Design, Code 311
San Diego, CA 92152

LCDR Michael O'Bar
Navy Personnel Research and
  Development Center
Code 305
San Diego, CA 92152

Naval Training Equipment Center
ATTN: Technical Library
Orlando, FL 32813

Human Factors Department
Code N215
Naval Training Equipment Center
Orlando, Fl 32813

Dr. Alfred F. Smode
Training Analysis and Evaluation Group
Naval Training Equipment Center
Code N-OOT
Orlando, Fl 32813

Dr. Gary Poock
Operations Research Department
Naval Postgraduate School
Monterey, CA 93940

Mr. J. Barber
Headquarters, Department of the
  Army, DAPE-PBR
Washington, D.C. 20546

Dr. Joseph Zeidner
Director, Organization and
  Systems Research Laboratory
U.S. Army Research Institute
1300 Wilson Boulevard
Arlington, VA 22209

Dr. Edgar M. Johnson
Organization and Systems
  Research Laboratory
U.S. Army Research Lab
1300 Wilson Blvd
Arlington, VA 22209

ARI Field Unit - Leavenworth
Post Office Box 3122
Fort Leavenworth, Kansas 66027

U.S. Air Force Office of Scientific
  Research
Life Sciences Directorate, NL
Bolling Air Force Base
Washington, D.C. 20332

Dr. Donald A. Topmiller
Chief, Systems Engineering Branch
Human Engineering Division
USAF AMRL/HES
Wright-Patterson AFB, OH 45433

Lt. Col. Joseph A. Birt
Human Engineering Division
Aerospace Medical Research Laboratory
Wright Patterson AFB, OH 45433

Air University Library
Maxwell Air Force Base, AL 36112

Dr. Robert Williges
Human Factors Laboratory
Virginia Polytechnic Institute
130 Whittemore Hall
Blacksburg, VA 24061

Dr. Arthur I. Siegel
Applied Psychological Services, Inc.
404 East Lancaster Street
Wayne, PA 19087

Dr. Robert R. Mackie
Human Factors Research, Inc.
Santa Barbara Research Park
6780 Cortona Drive
Goleta, CA 93017

Dr. Gershon Weltman
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, CA 91364

Dr. Paul Slovic
Decision Research
1201 Oak Street
Eugene, OR 97410

Dr. Terence R. Mitchell
University of Washington
Seattle, WA 98195

Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army-Navy Drive
Arlington, VA 22202

Journal Supplement Abstract Service
American Psychological Association
1200 17th Street, N.W.
Washington, D.c. 20036 (6 cys)

Dr. William A. McClelland
Human Resources Research Office
300 N. Washington Street
Alexandria, VA 22314

Director, Human Factors Wing
Defense & Civil Institute of
Environmental Medicine
Post Office Box 2000
Downsville, Toronto, Ontario
CANADA

Dr. A.D. Baddeley
Director, Applied Psychology Unit
Medical Research Council
15 Chaucer Road
Cambridge, CB2 2EF
ENGLAND

Dr. David Zaidel
Road Safety Centre Technion City
Haifa
ISRAEL